arXiv:1402.4884v2 [stat.ML] 21 Feb 2014

Le Cam meets LeCun: Deficiency and Generic Feature Learning

Brendan van Rooyen BRENDAN.VANROOYEN@ANU.EDU.AU
Robert C. Williamson BOB.WILLIAMSON@ANU.EDU.AU
Australian National University and NICTA, Canberra ACT 0200, Australia

Abstract

"Deep Learning" methods attempt to learn generic features in an unsupervised fashion from a large unlabelled data set. These generic features should perform as well as the best hand crafted features for any learning problem that makes use of this data. We provide a definition of generic features, characterize when it is possible to learn them and provide algorithms closely related to the deep belief network and autoencoders of deep learning. In order to do so we use the notion of deficiency distance and illustrate its value in studying certain general learning problems.

1. Introduction

"Deep" unsupervised feature learning methods (Bengio, 2009; Hinton and Salakhutdinov, 2006; Vincent et al., 2008; Tenenbaum et al., 2000; LeCun, 2013) present a challenge to learning theory. This paper takes up this challenge, of explaining when and why these techniques work. Is it possible to learn generic features from data in an unsupervised fashion that perform well in a multitude of tasks, and if so how do we learn these features? Following the work of the statistician Lucien Le Cam (Lecam, 2011, 1964, 1974) and utilizing the techniques of statistical decision theory, in particular the comparison of statistical experiments (Blackwell, 1951, 1953; Torgersen, 1991; Ferguson, 1967) we show (theorem 3) that it is possible to construct generic features Z from data X if and only if one can find a encoder/decoder pair

$$X \xrightarrow{\text{encoder}} Z \xrightarrow{\text{decoder}} X$$

with low probability of reconstruction error. Furthermore, the worst case difference in performance of the best decision rule that uses such features versus the best decision rule that uses the raw data is bounded above by the probability of reconstruction error. We also show that we can learn this encoder/ decoder pairing in a hierarchical fashion and that the probability of reconstruction error of such a"stacked" system is bounded by the sum of the probability of reconstruction errors of each layer.

While our approach is abstract, the ultimate pay off will be a novel inequality (theorem 2) that provides a characterization of when generic feature learning is possible. This inequality coupled with the concept of deficiency (Lecam, 2011; Torgersen, 1991) (to be explained in the paper) illuminates the algorithms used in deep learning and provides means to judge the generic quality of the features learnt by such methods.

2. The General Learning Problem

For all of the following assume that all of the measure spaces Θ, A, X, Y, Z and so on are finite. This does not restrict any of the results, rather it allows for a cleaner presentation free of measure theoretic technicalities as well as boundedness and existence concerns.

A learning problem is a quintuple (Θ, X, T, A, L). Θ is a set of possible "true hypothesis" or unknowns. While we cannot observe Θ directly, we can observe data in some set X. T is a relationship between the two sets Θ and X called the *experiment*. $T(\theta)$ tells us what data we expect to see if θ is the true hypotheses. Ultimately we are required to make a decision by choosing an action $a \in A$, and our performance is measured by a loss function $L : \Theta \times A \to \mathbb{R}$. We view the loss as an integral part of a learning problem and as such do not place any restrictions on it other than boundedness. As is usual in statistical learning, for our possible relationships we use markov kernels (conditional probability assignment/stochastic matrices).

Definition 1 *A* Markov kernel $T : \Theta \rightsquigarrow X$ *is a function from* Θ *to* $\mathcal{P}(X)$*, the set of probability distributions on* X.

For arbitrary sets Y and Z, denote by $M(Y, Z)$, the set of all Markov kernels from Y to Z. As we can represent $\mathcal{P}(Y)$ by vectors in $\mathbb{R}^{|Y|}$ with positive entries we have $\mathcal{P}(Y) \subset \mathbb{R}^{|Y|}$. As such one can represent a Markov kernel $T : Y \rightsquigarrow Z$ as an $|Z| \times |Y|$ matrix of positive entries where the sum of all entries in each column is equal to 1. It is easily verified that $M(Y, Z)$ is a closed convex subset of $\mathbb{R}^{|Z|\times|Y|}$, the set of all $|Z| \times |Y|$ matrices.

A function $f : Y \to Z$ defines a Markov kernel F with $F(y) = \delta_{f(y)}$, a point mass distribution on $f(y)$. For every measure space X, there are two special Markov kernels, the identity (or completely informative) Markov kernel from the identity function $\mathrm{id}_X : X \to X$, and the completely uninformative Markov kernel from the function $\bullet_X : X \to \bullet$ from X to a one element set $\bullet$.

From a prior distribution $\pi \in \mathcal{P}(\Theta)$ and a Markov kernel $T : \Theta \rightsquigarrow X$ we can construct a joint distribution $T \otimes \pi \in \mathcal{P}(\Theta \times X)$. Using the matrix vector representation, this is achieved by post multiplying T with a diagonal matrix with π on the diagonal, $T \otimes \pi = T \operatorname{diag}(\pi)$. This is no different to the standard product rule $P(\theta, x) = P(x|\theta)P(\theta)$.

Given a prior distribution $\pi \in \mathcal{P}(\Theta)$ and a Markov kernel $T : \Theta \rightsquigarrow X$ we denote the Markov kernel obtained by Bayes rule by $T^* : X \rightsquigarrow \Theta$.

A learning problem can more compactly be represented as the pair (L, T) where Θ, A, X can be inferred from the type signatures of L and T. We measure the size of loss functions by $\|L\|_\infty = \sup_{\theta,a} |L(\theta, a)|$

2.1. Decision Rules

Upon observing data $x \in X$ we are required to relate x to a set of actions A by some other Markov kernel $d : X \rightsquigarrow A$ known as a (randomized) decision rule.

$$\Theta \overset{T}{\rightsquigarrow} X \overset{d}{\rightsquigarrow} A$$

We are judged on the quality of the composed relation $d \circ T : \Theta \rightsquigarrow A$.

Definition (Composition) *Suppose $T_1 : X \rightsquigarrow Y$ and $T_2 : Y \rightsquigarrow Z$ are Markov kernels. Then we can compose T_1 and T_2 yielding $T_3 : X \rightsquigarrow Z$ by matrix multiplication*

$$T_3 = T_2 \circ T_1 = T_2 T_1.$$

A Markov kernel $T : X \rightsquigarrow Y$ provides a function $T : \mathcal{P}(X) \to \mathcal{P}(Y)$ by matrix multiplication. To calculate $T(\pi), \pi \in \mathcal{P}(X)$ we identify π with a vector and T with a matrix and use matrix multiplication. This function is convex linear.

2.2. Risk and Value: Ranking Decision Rules and Learning Problems

Given a learning problem (Θ, X, T, A, L) one can rank decision rules $d : X \rightsquigarrow A$ using the full Bayes risk

$$\begin{aligned} R_L : \mathcal{P}(\Theta) \times M(\Theta, A) &\to \mathbb{R} \\ (\pi, D) &\mapsto \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{a \sim D(\theta)} L(\theta, a). \end{aligned}$$

Here $\pi \in \mathcal{P}(\Theta)$ is a prior distribution on Θ which reflects which hypotheses we feel are more or less likely to be true. Note that both $\mathcal{P}(\Theta)$ and $M(\Theta, A)$ are convex sets and that R_L is convex bilinear (the same as bilinear but restricted to convex combinations). Alternately, taking a supremum over the prior yields the max risk.

Ranking Learning Problems. We also rank the difficulty of learning problems. The greatest challenge in a learning problem comes from the fact we can not use an arbitrary decision rule $D \in M(\Theta, A)$. Rather, we are restricted to a certain subset of $M(\Theta, A)$ that "factors through" T. We are only allowed to use the data we see.

Definition 2 (Factoring Through) *Suppose we have two Markov kernels $T : X \rightsquigarrow Z$ and $U : X \rightsquigarrow Y$. We say that U* factors through *T (written $T|U$) if there exists a Markov kernel $U/T : Z \rightsquigarrow Y$ such that $U = (U/T) \circ T$. Denote by*

$$M(X, Y)_T := \{U \in M(X, Y) : U = (U/T) \circ T \text{ for some } U/T \in M(Z, Y)\}.$$

If $T|U$ then U can be thought of as T with extra noise U/T. The reader is directed to section 1 of the appendix for more properties of factoring through. In this notation $d = D/T$. If $T|D$ we have

$$\begin{aligned} R_L(\pi, D) &= R_L(\pi, (D/T) \circ T) \\ &= \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{a \sim (D/T) \circ T(\theta)} L(\theta, a) \\ &= \mathbb{E}_{\theta \sim \pi} \mathbb{E}_{x \sim T(\theta)} \mathbb{E}_{a \sim (D/T)(x)} L(\theta, a). \end{aligned}$$

We assign a *value*

$$\mathcal{V}(\pi, \Theta, X, T, A, L) = \mathcal{V}_L(\pi, T) := \inf_{D \in M(\Theta, A)_T} R_L(\pi, D)$$

to a learning problem, with lower value being better. Taking a supremum of the value over the prior yields the minimax risk. The value is the risk of the best possible decision rule for the learning problem at hand.

Bayes Decision Rules are Optimal. If we use $V_L(\pi, T)$ to order decision rules then the best $D \in M(\Theta, A)_T$ is found by using Bayes rule

$$(D/T)(x) = \arg\inf_a \mathbb{E}_{\theta \sim T^*(x)} L(\theta, a),$$

with risk

$$R_L(\pi, D) = \mathcal{V}_L(\pi, T) = \mathbb{E}_{x \sim \pi_X} \inf_a \mathbb{E}_{\theta \sim T^*(x)} L(\theta, a) = \mathbb{E}_{x \sim \pi_X} \underline{L}(T^*(x)),$$

where $\underline{L} : \mathcal{P}(\Theta) \to \mathbb{R}$, $\underline{L}(\pi) = \inf_a \mathbb{E}_{\theta \sim \pi} L(\theta, a)$. Hence $\underline{L}$ is concave. This results allows us to parametrize the action set A by $\mathcal{P}(\Theta)$, by taking

$$\begin{aligned} f : \mathcal{P}(\Theta) &\to A \\ \pi &\mapsto \arg\inf_a \mathbb{E}_{\theta \sim P} L(\theta, a) \end{aligned}$$

effectively properising the loss function. $\hat{L}(\theta, \pi) := L(\theta, f(\pi))$ is a proper loss (Reid and Williamson, 2011; Dawid, 2007; Grünwald and Dawid, 2004; Parry et al., 2012). There are deep connections between $\underline{L}$ and $\hat{L}$, we review some of these in section two of the appendix.

Connections to other Information Measures.

There are many connections between $V_L(\pi, T)$ and different information measures present in the literature.

Definition *For a convex $f : \mathbb{R}_+^{n-1} \to \mathbb{R}$ the f-information of a set of n distributions $P_1, \ldots, P_n \in \mathcal{P}(X)$ is*

$$I_f(P_1, \ldots, P_n) := \int_X f(\frac{dP_2}{dP_1}, \ldots, \frac{dP_n}{dP_1}) dP_1.$$

f-informations are a multi distribution extension of f divergences and are used in certain generalizations of rate-distortion theory where they produce better bounds than the standard techniques (Ziv and Zakai, 1973; Zakai and Ziv, 1975; Reid and Williamson, 2011; Garcia-Garcia and Williamson, 2012). For suitable choices of f one can recover more known measures of information such as mutual information.

Theorem *For all experiments T, loss functions L and priors π, the gap between then the value of T and the least informative experiment $\bullet_\Theta$ is a f-information for suitable f*

$$\mathcal{V}_L(\pi, \bullet_\Theta) - \mathcal{V}_L(\pi, T) = I_f(T(\theta_1), \ldots, T(\theta_n)) = I_f(T)$$

We direct the reader to Reid and Williamson (2011); Garcia-Garcia and Williamson (2012). By the bijections presented in these two papers, one can replace $V_L(\pi, T)$ by these divergences in all that follows. In particular any $|\mathcal{V}_L(\pi, T) - \mathcal{V}_L(\pi, U)|$ can be replaced with a $|I_f(T) - I_f(U)|$ for suitable f, with no effect on the result. The reader is directed to section 3 of the appendix for proof.

2.3. Examples

Here we present some examples of familiar learning problems phrased in this more abstract language. Normally there is a distinction between *learning algorithms*, something that takes a data set of n instance-label pairs and produces a classifier, and a *decision rule* that is the learnt classifier. Here we do not make such a distinction. Both learning algorithms and decision rules produce *actions*, hence we only use the term decision rule.

Example 1 (Classification) *$\Theta = \{-1, 1\}$ and a Markov kernel $T : \Theta \rightsquigarrow X$ is then a pair of distributions $T(1) = P$, $T(-1) = Q$ on X. Normally $A = \Theta$ and a decision rule picks the corresponding label for a given observed $x \in X$. Different losses could be used, eg the 0-1 loss L_{01}.*

Example 2 (Supervised Learning) *There is a space of labels $Y = \{-1, 1\}$ and a space of covariates Z with $X = (Z \times Y)^n$. $\Theta = \mathcal{P}(Z \times Y)$ the space of joint distributions on $Z \times Y$ with $T : \Theta \rightsquigarrow X$ the map that sends each distribution to its n-fold product. A is then some set of classifiers (eg linear hyperplanes/kernel machines and so on). A decision rule then produces a classifier a from n pairs (z_i, y_i). For example empirical risk minimization algorithms, $ERM : X \to A$, pick the classifier that minimizes the empirical loss on the observed training set. Many suitable losses exists but normally $L(\theta, a)$ is the misclassification probability of the classifier a when used against the distribution θ.*

Example 3 (Generalized Supervised Learning) *Θ and A are the same as supervised learning, although the data observed is different. For example in semi-supervised learning we have $X = (Z \times Y)^n \times Z^m$, we observe n instance label pairs and m instances. $T : \Theta \rightsquigarrow X$ then maps each joint distribution θ to a product of n copies of itself and m copies of its marginal distribution over instances.*

Example 4 (Active Learning) *Θ could be anything with $X = Z^n$, length n sequences in some set Z. Each active learning policy determines a different $T : \Theta \rightsquigarrow X$.*

3. Feature Learning

Starting from a learning problem (Θ, X, T, A, L), feature learning methods aim to extract features $\phi : X \rightsquigarrow Z$. One then bases all decisions on these features.

$$\Theta \overset{T}{\rightsquigarrow} X \overset{\phi}{\rightsquigarrow} Z \overset{d}{\rightsquigarrow} A$$

These methods swap the original learning problem with $(\Theta, Z, \phi \circ T, A, L)$. Normally the space Z is smaller/ of lower dimension than X and aims at presenting a "compressed" view of the information contained in X. Features can be used for several reasons including communication/ storage constraints, increased performance (ie by implementing decision rules based on Z rather than X directly), knowledge discovery and to avoid "curse of dimensionality" problems.

3.1. Supervised Feature Learning

There has been much attention in the Machine Learning literature on supervised feature learning techniques, were Θ, T, L and the prior π are fixed. These methods construct features by minimizing the *feature gap*

$$\Delta\mathcal{V}_L(\pi, \phi, T) := \mathcal{V}_L(\pi, \phi \circ T) - \mathcal{V}_L(\pi, T).$$

There is now a general framework for solving such problems based largely on variations of the Blahut-Arimoto Algorithm from Rate Distortion theory (Banerjee et al., 2005; Tishby et al., 1999;

Cover and Thomas, 2012). For particular choices of T and L these methods reproduce many clustering methods such as k-means. We review these methods in section 4 of the appendix. These feature learning methods are not general enough for our purposes as they rely on both the experiment and loss. For example if we wish to *learn* T from data given by (θ, x) pairs, then we are required to learn the features after we have learnt the experiment T. Ideally we would like to learn a feature map $\phi : X \rightsquigarrow Z$ independently from T so that learning $\phi \circ T$ is just as beneficial as learning T no matter what T is.

3.2. Generic Feature Learning

For many learning problems, a large amount of unlabelled/loosely labelled data X is readily available. For example, with any problem involving images one only has to enter some basic search queries into google to be presented with millions of instances. One of the main arguments of the deep learning community is that while this data may not be of direct use in learning classifiers, it can be of great use in learning feature representations. There exists many methods in the literature to learn features from unlabelled data.

In line with these methods we consider the following relaxation of the supervised feature learning problem. We assume that Θ, T, A, L and the prior π are allowed to vary leaving only X fixed, with one restriction. We assume that there is enough unlabelled data collected from the marginal distribution π_X on X that we are able to form an accurate estimate of π_X. We consider all learning problems and priors π that are consistent with this information about X ie with $T(\pi) = \pi_X$. We then seek to find features ϕ so that the value of $\mathcal{V}_L(\pi, \phi \circ T)$ is as close to $\mathcal{V}_L(\pi, T)$ as possible, no matter what what π, Θ, T, A, L are. To ensure that minor differences in value are not exploited by multiplying the loss function by a large constant, we penalize the size of the loss function by $\|L\|_\infty$.

Definition 3 (Generic Features) *Fix a measure space X and a distribution π_X. $\phi : X \rightsquigarrow Z$ are* generic features of quality ϵ *for X if for all learning problems (Θ, X, T, A, L) and priors π with $T(\pi) = \pi_X$ we have*

$$\Delta\mathcal{V}_L(\pi, \phi, T) \leq \epsilon \|L\|_\infty.$$

Ideally we want to make ϵ as small as possible, and if ϵ is 0 then our features do not ever decrease the value.

For our more relaxed problem the value of our features is effectively

$$\begin{aligned} &\sup_{L:\|L\|_\infty \leq 1} \sup_{T} \sup_{\pi\,:\,T(\pi)=\pi_X} \Delta\mathcal{V}_L(\pi, \phi, T) \\ &= \sup_{\Theta} \sup_{A} \sup_{L \in \mathbb{R}^{\Theta \times A} : \|L\|_\infty \leq 1} \sup_{T} \sup_{\pi\,:\,T(\pi)=\pi_X} \Delta\mathcal{V}(\Theta, A, X, \phi, T, \pi) \end{aligned}$$

Luckily supremums like these have been tackled in theoretical statistics particularly in the work of Lucien Le Cam (Lecam, 2011, 1964, 1974). In his 1964 paper Le Cam coined the deficiency distance as an extension of David Blackwell's ordering of experiments Blackwell (1953, 1951) and as a means to provide an approximate version of the statistical notion of sufficiency. This quantity was used in his later work to form a metric not just on probability distributions but on experiments, and in particular allows one to calculate supremums over all loss functions and priors with fixed Θ, T. We introduce this quantity (the deficiency) in the next section.

4. Approximate Factoring Through and Deficiency

Suppose $T : \Theta \rightsquigarrow X$ and $U : \Theta \rightsquigarrow Y$ are Markov kernels where U does not factor through T. We measure the degree to which U fails to factor through T by the *weighted directed deficiency* (Torgersen, 1991).

$$\delta_\pi(T, U) := \inf_{V:X\rightsquigarrow Y} \mathbb{E}_{\theta\sim\pi}\|U(\theta) - V \circ T(\theta)\|.$$

$\|P - Q\|$ is the variational divergence between the distributions $P, Q \in \mathcal{P}(X)$, a standard metric on probability distributions (see section 5 of the appendix for properties). Calculating weighted directed deficiencies is a convex (actually linear) optimization problem. One has

$$\begin{aligned} f(\pi, V) &= \mathbb{E}_{\theta\sim\pi}\|U(\theta) - V \circ T(\theta)\| \\ &= \mathbb{E}_{\theta\sim\pi}\|U(\theta) - V \circ T(\theta)\|_1 \\ &= \|U \otimes \pi - (V \circ T) \otimes \pi\|_1 \\ &= \|U \ \mathrm{diag}(\pi) - V \ T \ \mathrm{diag}(\pi)\|_1 \end{aligned}$$

is linear in π. Since the variational divergence $||P - Q||$ is convex in Q, f is also convex in V, because $\|U \ \mathrm{diag}(\pi) - V \ T \ \mathrm{diag}(\pi)\|_1$ is the composition of a linear function and a convex function. Hence determining weighted directed deficiencies is a l_1 minimization problem. Fast methods exist for solving this problem (eg the well known simplex method of linear programming).

Taking a supremum over the prior π yields the *directed deficiency*,

$$\delta(T, U) := \sup_\pi \delta_\pi(T, U) = \inf_{V:X\rightsquigarrow Y} \sup_\theta \|U(\theta) - V \circ T(\theta)\|.$$

where the second follows from the minimax theorem (Komiya, 1988). For the sake of checking whether $T|U$ it suffices to use the weighted directed deficiency and a prior that does not put zero probability on any θ. In this case $\delta_\pi(T, U) = 0$ if and only if $T|U$ (Torgersen, 1991). The *weighted deficiency*, and *deficiency* are respectively

$$\begin{aligned} \Delta_\pi(T, U) &:= \max(\delta_\pi(T, U), \delta_\pi(U, T)) \\ \Delta(T, U) &:= \sup_\pi \Delta_\pi(U, T). \end{aligned}$$

$\Delta(T, U) = 0$ if and only if $T|U$ and $U|T$, when T is *isomorphic* to U written $T \cong U$. The deficiency distance is a true metric on the space of experiments (modulo isomorphic experiments). A proof of this is included in the appendix.

4.1. Relation to Risk

Factoring through and approximate factoring through are deeply related to the worst case difference in performance between two learning problems with the same Θ as the loss is varied. Here we state the three theorems that highlight the connections between factoring through and risk (Torgersen, 1991). Fix Θ and two experiments $T : \Theta \rightsquigarrow X$ and $U : \Theta \rightsquigarrow Y$.

Theorem (Information Processing) *If $T|U$ then for any loss function L and prior π* $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U)$.

In particular the information processing theorem implies that $\Delta\mathcal{V}_L(\pi, \phi, T) \geq 0$.

Theorem (Blackwell-Sherman-Stein) *$T|U$ if and only if $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U)$ for all loss functions L and priors π.*

Theorem (Randomization) *Fix $\epsilon > 0$, T, U and π. $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U) + \epsilon\|L\|_\infty$ if and only if $\delta_\pi(T, U) \leq \epsilon$ for all loss functions L.*

These three theorems allow one to move between decision theoretic notions such as risk and value to probability theoretic notions such as factoring through. For example the original definition of sufficiency can be interpreted in terms of factoring through.

Theorem *Fix an experiment $T : \Theta \rightsquigarrow X$ and a function $f : X \to Y$. f is a* sufficient statistic *if $f \circ T \cong T$.*

By the Blackwell-Sherman-Stein theorem we have an equivalent condition for sufficiency in terms of value.

Theorem *Fix an experiment $T : \Theta \rightsquigarrow X$ and a function $f : X \to Y$.Then f is a sufficient statistic if $\mathcal{V}_L(\pi, f \circ T) = \mathcal{V}_L(\pi, f \circ T)$ for all L and π.*

Isomorphic experiments always have the same value, no matter what the loss function or the set of actions. Approximately isomorphic experiments, ones where $\Delta(T, U)$ is small, always have approximately the same value. Due to the similarities between learning features and sufficiency statistics, it should be of no surprise that tools for working with approximate sufficiency appear in feature learning.

The Randomization theorem is an example of an approximate notion in probability theory (here approximate sufficiency) has a dual approximate notion in terms of risk.

Theorem 1 *For all experiments U, T and all priors π*

$$\Delta_\pi(U, T) = \sup_L \frac{|\mathcal{V}_L(\pi, U) - \mathcal{V}_L(\pi, T)|}{\|L\|_\infty}$$

This result is an improvement and generalization of the result contained in Liese (2012) that applied only to binary experiments, $\Theta \cong \{-1, 1\}$, and held with inequality. The proof is included in the appendix. We utilize this theorem and the randomisation theorem heavily in the following section.

4.2. Reductions via Factoring Through

Approximate factoring through can be used to transform decision rules for one learning problem to rules for another, with a provable bound on the performance of this decision rule. Suppose $D_U \in M(\Theta, A)_U$ is a decision rule used for the learning problem (Θ, Y, U, A, L). For another learning problem (Θ, X, T, A, L), we can construct a decision rule from a Markov kernel $V : X \rightsquigarrow Y$

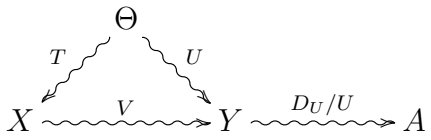

$D_T = (D_U/U) \circ V \circ T$. Furthermore if $\epsilon = \mathbb{E}_{\theta\sim\pi}\|U(\theta) - V \circ T(\theta)\|$ then

$$R_L(\pi, D_T) \leq R_L(\pi, D_U) + \epsilon\|L\|_\infty.$$

By taking infimums over V we obtain the smallest ϵ in the above.

5. Analysis of Generic Feature Learning via Deficiency

Assume one has enough data in some measure space X to form a good estimate of the marginal distribution π_X. We wish to construct generic features $\phi : X \rightsquigarrow Z$ for X. This is equivalent to finding a ϕ that minimizes $\sup_T \Delta_\pi(T, \phi \circ T)$. One might imagine that this means finding for each ϕ the worst T, but this is not the case.

Theorem 2 *For all experiments $T : \Theta \rightsquigarrow X$, for all measure spaces Z and for all feature maps $\phi : X \rightsquigarrow Z$*

$$\Delta_\pi(T, \phi \circ T) \leq \Delta_{T(\pi)}(\mathrm{id}_X, \phi).$$

No matter which feature map ϕ we use, the worst learning problem we can pit against it is the one that asks you to reconstruct X directly from the features. The proof is straightforward and is included in the appendix. It hinges on the representation of $\mathcal{V}_L(\pi, T)$ in terms of average posterior Bayes risk and the Randomization theorem. By definition and as $\mathrm{id}_X \,|\phi$

$$\begin{aligned}
\Delta_{\pi_X}(\mathrm{id}_X, \phi) &= \delta_{\pi_X}(\phi, \mathrm{id}_X) \\
&= \inf_{d: Z \rightsquigarrow X} \mathbb{E}_{x \sim \pi_X} \|\mathrm{id}_X(x) - (d \circ \phi)(x)\| \\
&= \inf_{d: Z \rightsquigarrow X} \mathbb{E}_{x \sim \pi_X} \|\delta_x - (d \circ \phi)(x)\| \qquad (1) \\
&= 2 \inf_{d: Z \rightsquigarrow X} \mathbb{E}_{x \sim \pi_X} \mathbb{E}_{x' \sim (d \circ \phi)(x)} \mathbb{1}(x' = x), \qquad (2)
\end{aligned}$$

where from lines (1) to (2) we have used one of the equivalent forms of the variational divergence listed in the appendix. Hence $\Delta_{\pi_X}(\mathrm{id}_X, \phi)$ is equal to twice the minimal possible average reconstruction error from the encoder ϕ and the prior π_X. This means that finding the best generic features for the data X involves finding the ϕ that gives the learning problem (X, Z, ϕ, X, L_{01}) the highest value $\mathcal{V}_{L_{01}}(\pi_X, \phi)$. We term this problem the *reconstruction problem*.

Theorem 3 *Fix a prior π_X and an $\epsilon > 0$. ϕ constitutes generic features of quality ϵ for X if and only if for the reconstruction problem (X, Z, ϕ, X, L_{01}), one can find a decision rule $d : Z \rightsquigarrow X$ with*

$$R_{L_{01}}(\pi_X, d \circ \phi) \leq \frac{\epsilon}{2}.$$

For a proof see appendix. If we optimize over both the encoder ϕ and the decoder d, finding

$$\inf_{\phi, d} \mathbb{E}_{x \sim \pi_X} \|\mathrm{id}_X(x) - (d \circ \phi)(x)\|$$

we obtain a variant of the popular autoencoder algorithm from deep learning (Vincent et al., 2008).

Of course one can always take $\phi = \mathrm{id}_X$ in which case no real feature learning is done and no performance is lost. However, it is more instructive to set Z to a measure space of smaller size/ dimension than X so that the feature learning extracts (provably) useful patterns in π_X. Performing the joint minimization is a non-convex problem.

5.1. Relation to other Feature Learning Methods

Infomax. When faced with a choice of feature maps $\phi_i : X \rightsquigarrow Z_i$, the Infomax principle (Bell and Sejnowski, 1995; Linsker, 1989) dictates that you should choose the features that minimize the conditional entropy between data and features, $H(X|Z_i)$. By the Hellman-Raviv inequality from information theory (Hellman and Raviv, 1970) we have

$$\Delta_{\pi_X}(\mathrm{id}_X, \phi_i) \leq H(X|Z_i)$$

meaning the Infomax principle is minimizing an upper bound of the reconstruction error.

Manifold Learning. Under the assumption that π_X has support on some manifold $M \subseteq X$ manifold learning methods aim to extract this manifold and provide a parametrization $\phi : M \to \mathbb{R}^n$ (Silva and Tenenbaum, 2002; Belkin and Niyogi, 2003). If we are able to learn this manifold then any coordinate system would constitute generic features.

Sparse Coding. Much like manifold learning, sparse coding also attempts to find lower dimensional structure in π_X (Lee et al., 2006; Olshausen and Field, 1997). Here Z is chosen to have higher dimension than X, however image of the feature map $\phi : X \rightsquigarrow Z$ should comprise only of sparse vectors, those with few non-zero entries. If ϕ is injective on the support of X then ϕ are generic features.

5.2. Learning Feature Hierarchies

One of the tenets of the deep learning paradigm is that features should be learnt in a hierarchical fashion. One should first find patterns in π_X through a feature map ϕ and then find patterns in $\phi(\pi_X)$ and so on. We construct a chain of feature maps

$$X \overset{\phi_1}{\rightsquigarrow} Z_1 \overset{\phi_2}{\rightsquigarrow} Z_2 \overset{\phi_3}{\rightsquigarrow} \dots$$

with final feature space Z_n and final feature map given by the composition of all maps in the chain. Proceeding in this fashion allows greater control over the feature spaces Z_i. For example the first could be of similar (but still lower) size than X, perhaps before a big drop off in the middle of the chain. If we can learn each of the ϕ_i iteratively it also makes searching for features easier. For example, perhaps the first three feature maps have low probability of reconstruction error but not the fourth. In this situation at least we still have a good feature map given by the composition of the first three mappings. We have for any chain of feature maps

$$\begin{aligned}
&\Delta_{\pi_X}(\mathrm{id}_X, \phi_n \circ \cdots \circ \phi_2 \circ \phi_1) \\
\leq &\Delta_{\pi_X}(\mathrm{id}_X, \phi_1) + \Delta_{\pi_X}(\phi_1, \phi_2 \circ \phi_1) + \cdots + \Delta_{\pi_X}(\phi_{n-1} \circ \cdots \circ \phi_1, \phi_n \circ \cdots \circ \phi_1) \qquad (1) \\
\leq &\Delta_{\pi_X}(\mathrm{id}_X, \phi_1) + \Delta_{\phi_1(\pi_X)}(\mathrm{id}_{Z_1}, \phi_2) + \cdots + \Delta_{\phi_{n-1} \circ \cdots \circ \phi_2 \circ \phi_1(\pi_X)}(\mathrm{id}_{Z_{n-1}}, \phi_n) \qquad (2)
\end{aligned}$$

where (1) follows as the weighted deficiency satisfies a triangle inequality and (2) follows from repeated application of theorem 2. The reconstruction error of the entire system is bounded by the sum of reconstruction errors at each step of the chain. This means we can learn a feature mapping iteratively, by first learning patterns in π_X, then in $\phi(\pi_X)$ and so on. This is exactly the process that occurs in a Deep Belief Network (Hinton and Salakhutdinov, 2006)

5.3. Supervised Feature learning can work when Generic Feature Learning Fails

We present two examples where one can not learn generic features, however we can learn experiment/loss specific features.

Experiment Specific Features. Let $\Theta = \mathbb{R}$ with $X = \mathbb{R}^n$ and $T(\theta)$ given by the product of n normal distributions with mean θ and variance 1. It is easy to verify that the sample mean $\phi : X \to \mathbb{R}$ is a sufficient statistic meaning that at least for this experiment we can greatly compress the information contained in X. However, if we take as a prior π for Θ a normal distribution of mean 0 and variance 1, then the marginal distribution π_X will not be concentrated on a set of smaller dimension nor have any particularly interesting structure. Hence we can not find interesting generic features in this case.

Experiment and Loss Specific Features. Let $\Theta = \{-1, 1\}$ with $T(\theta)$ a normal distribution centred on θ as in the figure below.

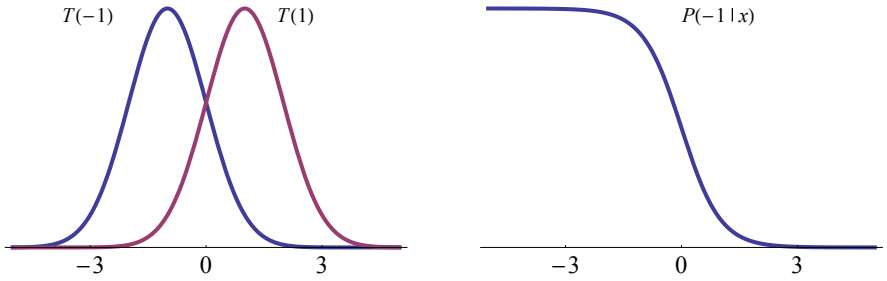

Figure 1: Figures for Loss Specific Feature Learning, see text

For this experiment ,L_{01} and a uniform prior π, the best decision $d(x) = 1$ if $x > 0$ as $P(-1|x) > \frac{1}{2}$ and $d(x) = -1$ otherwise as $P(-1|x) \leq \frac{1}{2}$. It is easy to show that $\Delta\mathcal{V}_{L_{01}}(\pi, d, T) = 0$, all we need is the output of d. However if we change the loss to a cost sensitive loss L_c where say misclassifying a 1 is more costly than a -1, we no longer have $\Delta\mathcal{V}_{L_c}(\pi, d, T) = 0$.

5.4. Alternate Reconstruction Problems

The reconstruction problem (X, Z, ϕ, X, L_{01}) that is required to be solved to construct generic features is the most difficult one we can pose. To perform well in this problem we are required to reconstruct each $x \in X$ *exactly*. This discards other interesting structure the set X may have. For example if X is image data a different loss function $L : X \times X \to \mathbb{R}$, perhaps one elicited from psychological tests of what humans perceive to be different images is more appropriate. While these are valid points, we remind the reader that this is a first step in understanding these methods, and making any extra assumptions about X and its structure is exactly what we are trying to avoid. However, the Hellman Raviv inequality does give means of bounding the value $\mathcal{V}_{L_{01}}(\pi, \phi)$ with the value of different reconstruction problems.

6. Concluding Remarks and Future Work

We have defined generic features and have provided a characterization of when it is possible to learn them. In doing so we have illuminated some popular feature learning methods including autoencoders, deep belief networks and the Infomax principle.. We have moved from supervised feature learning methods

$$\inf_{\phi : X \rightsquigarrow Z} \Delta \mathcal{V}_L(\pi, \phi, T)$$

where Θ, A, L, π are all *fixed* to

$$\inf_{\phi : X \rightsquigarrow Z} \sup_{L : \|L\|_\infty \leq 1} \sup_{T} \sup_{\pi \,:\, T(\pi) = \pi_X} \Delta \mathcal{V}_L(\pi, \phi, T) \tag{$\star$}$$

with almost *nothing fixed*. Equation $(\star)$ shows how difficult and general finding generic features is. It is reasonable to argue that in practice on does not require features that work for *all* experiments, *all* losses and *all* priors, which by existing results implies a quantification over *all* f-informations. This begs the question of *which* experiments, loss functions and priors to consider. We might not require *all* the information in X to be maintained, just enough to suit our purposes. This is analogous to the the problem of formalizing the notion of how much information is contained in an experiment. As argued long ago by Morris Degroot (DeGroot, 1962), even if one is doing an experiment to "gain information", eventually one *does something* with this "information" by choosing how to act, the consequence of which will be measured by some loss. Hence a more general theory of feature learning needs to be able to control the sensitivity to the loss, allowing one to move from supervised feature learning to generic feature learning.

A starting point would be to take fixed L and T lying in some subset of $M(\Theta, X)$ as occurs in robust statistics (Huber, 2011), or to allowing small perturbations in the loss. Deficiency can possibly play a role in the development of algorithms to learn features when we take these restricted supremums. There is scope to develop new quantities and theorems analogous to those for deficiency where instead of a supremum over all losses, one takes a supremum over some restricted subset. At present this is an open and uncharted area of both machine learning and theoretical statistics.

7.

7.1. More Properties of Factoring Through

Lemma *Factoring through has the following properties. For all sets W, X, Y, Z and Markov kernels $T_1 : X \rightsquigarrow Y$ and $T_2 : Y \rightsquigarrow W$ we have*

1. $M(X,Z)_{T_1} \supseteq M(X,Z)_{T_2 \circ T_1}$
2. $M(X,Z) = M(X,Z)_{id_X} \supseteq M(X,Z)_{T_1}$
3. $M(X,Z)_{T_1} \supseteq M(X,Z)_{\bullet_X}$

Proof For (1) if $D \in M(X,Z)_{T_2 \circ T_1}$ then we have

$$D = (D/(T_2 \circ T_1)) \circ (T_2 \circ T_1) = ((D/(T_2 \circ T_1)) \circ T_2) \circ T_1$$

which is obviously in $M(X,Z)_{T_1}$. For (2) note that for any Markov kernel $D : X \rightsquigarrow Y$ one has $D \circ id_X = D$. Hence all kernels factor through the identity. For (3) take any $D \in M(X,Z)_{\bullet_X}$ and recall that $\bullet_X = \bullet_Y \circ T_1$. Hence

$$D = (D/\bullet_X) \circ \bullet_X = (D/\bullet_X) \circ \bullet_Y \circ T_1$$

which is obviously in $M(X,Z)_{T_1}$ ■

Note that $M(X,Z)_{\bullet_X}$ comprises the constant Markov kernels, ones that map each $x \in X$ to the same distribution on Z.

One can view factoring through as adding noise. By showing that $T|U$, we are showing that U is T composed with extra noise U/T

For some intuition on what factoring through looks like below is a plot of four binary experiments ($\Theta = -1, 1$).

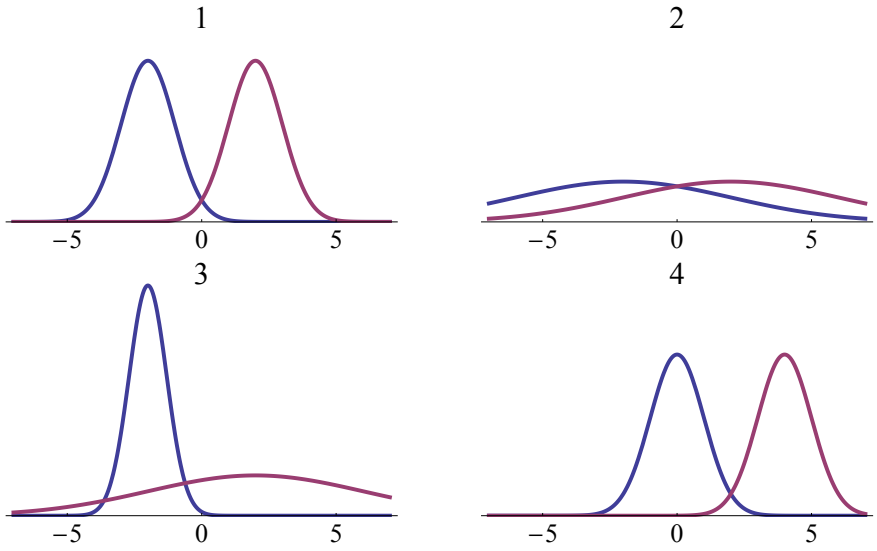

We have that:

- The second factors through the first.
- The third does not factor through the first (nor the first through the third).
- The fourth factors through the first and vice versa (it is just a shifted version of the first).

Suppose $T : \Theta \rightsquigarrow X$ and $U : \Theta \rightsquigarrow Y$ are two Markov kernels. If $T|U$ and $U|T$ then we say that T is isomorphic to U written $T \cong U$. Isomorphic Markov kernels can appear quite different. For example suppose that $T : \Theta \rightsquigarrow X$ is an exponential family distribution (Θ in this case are the parameters for the family). If $\phi : X \to \mathbb{R}^n$ is the sufficient statistics for the family T then $T \cong \phi \circ T$ even though they appear quite different, and that it appears ϕ may throw away lots of information. From a statistical point of view isomorphic Markov kernels are the same.

7.2. f-Information and Value

Theorem *For all experiments T, loss functions L and priors π, the gap between then the value of T and the least informative experiment $\bullet_\Theta$ is a f-information for suitable f*

$$\mathcal{V}_L(\pi, \bullet_\Theta) - \mathcal{V}_L(\pi, T) = I_f(T(\theta_1), \ldots, T(\theta_n)) = I_f(T)$$

We repeat the proof presented in Garcia-Garcia and Williamson (2012), which is a variant of one presented in DeGroot (1962).

Proof

$$\mathcal{V}_L(\pi, T) = \mathbb{E}_{x \sim T(\pi)} \underline{L}(T^*(x))$$

and

$$\mathcal{V}_L(\pi, \bullet_\Theta) = \underline{L}(\pi)$$

where $\underline{L} : \mathcal{P}(\Theta) \to \mathbb{R}$ is concave. Let $|\Theta| = n$, $\pi_i = \pi(\theta_i)$, $P_i = T(\theta_i)$ and $M = \sum_{i=1}^{n} \pi_i P_i$. Then $T^*(x) = (\pi_1 \frac{dP_1}{dM}(x), \pi_2 \frac{dP_2}{dM}(x), \ldots, \pi_n \frac{dP_n}{dM}(x))$ and

$$\begin{aligned}
\mathcal{V}_L(\pi, T) &= \int_X \underline{L}(\pi_1 \frac{dP_1}{dM}, \pi_2 \frac{dP_2}{dM}, \ldots, \pi_n \frac{dP_n}{dM}) dM \\
&= \int_X \underline{L}(\frac{dP_1}{dM}(\pi_1, \pi_2 \frac{dP_2}{dP_1}, \ldots, \pi_n \frac{dP_n}{dP_1})) \frac{dM}{dP_1} dP_1 \\
&= \int_X \underbrace{\underline{L}(\frac{1}{\pi_1 + \sum_{i=2}^{n} \pi_i \frac{dP_i}{dP_1}}(\pi_1, \pi_2 \frac{dP_2}{dP_1}, \ldots, \pi_n \frac{dP_n}{dP_1}))(\pi_1 + \sum_{i=2}^{n} \pi_i \frac{dP_i}{dP_1})}_{\phi(\frac{dP_2}{dP_1}, \ldots, \frac{dP_n}{dP_1})} \, dP_1
\end{aligned}$$

We need to show that $\phi : \mathbb{R}_+^{n-1} \to \mathbb{R}$ is concave. Note that $\phi = g \circ h$ where

$$\begin{aligned}
h : \mathbb{R}_+^{n-1} &\to \mathbb{R}_+^n \\
v &\mapsto \operatorname{diag}(\pi)(1, v)
\end{aligned}$$

prepends 1 to v and and multiplies by the prior, and

$$\begin{aligned} g : \mathbb{R}_+^n &\to \mathbb{R} \\ v &\mapsto \|v\|_1 \underline{L}(\frac{v}{\|v\|_1}). \end{aligned}$$

As h is affine, ϕ is concave if g is. For all $\lambda \in [0,1]$ and $v_1, v_2 \in \mathbb{R}_+^n$

$$\begin{aligned} g(\lambda v_1 + (1-\lambda)v_2) &= \|\lambda v_1 + (1-\lambda)v_2\|_1 \underline{L}(\frac{\lambda v_1 + (1-\lambda)v_2}{\|\lambda v_1 + (1-\lambda)v_2\|_1}) \\ &= \|\lambda v_1 + (1-\lambda)v_2\|_1 \underline{L}(\frac{\lambda\|v_1\|_1}{\|\lambda v_1 + (1-\lambda)v_2\|_1}\frac{v_1}{\|v_1\|_1} + \frac{(1-\lambda)\|v_2\|_1}{\|\lambda v_1 + (1-\lambda)v_2\|_1}\frac{v_2}{\|v_2\|_1}) \\ &\leq \lambda\|v_1\|_1 \underline{L}(\frac{v_1}{\|v_1\|_1}) + (1-\lambda)\|v_2\|_1 \underline{L}(\frac{v_2}{\|v_2\|_1}) \\ &= \lambda g(v_1) + (1-\lambda) g(v_2) \end{aligned}$$

where we have used the concavity of $\underline{L}$. Therefore ϕ is concave. Finally

$$\mathcal{V}_L(\pi, \bullet_\Theta) - \mathcal{V}_L(\pi, T) = \int_X \underline{L}(\pi) - \phi(\frac{dP_2}{dP_1}, \ldots, \frac{dP_n}{dP_1}) dP_1 = I_f(\frac{dP_2}{dP_1}, \ldots, \frac{dP_n}{dP_1})$$

with $f = \underline{L}(\pi) - \phi$ convex. This completes the proof.

■

The converse is also true. The proof follows by similar manipulations and an appeal to the properties of proper losses (see Reid and Williamson (2011); Garcia-Garcia and Williamson (2012)), and is not inlcuded.

7.3. Proper Loss Functions

Here we review material relating to the construction of proper loss functions (Reid and Williamson, 2011; Dawid, 2007; Parry et al., 2012; Grünwald and Dawid, 2004).

Definition *A loss function* $L : \Theta \times \mathcal{P}(\Theta) \to \mathbb{R}$ *is* proper *if for all* $P \in \mathcal{P}(\theta)$

$$P \in \underset{Q \in \mathcal{P}(\Theta)}{\arg\inf} \, \mathbb{E}_{\theta \sim P} L(\theta, Q)$$

Any loss function can be *properized*.

Theorem *Let* $L : \Theta \times A \to \mathbb{R}$ *be a loss. Define*

$$\begin{aligned} f : \mathcal{P}(\Theta) &\to A \\ P &\mapsto \underset{a}{\arg\inf} \, \mathbb{E}_{\theta \sim P} L(\theta, a) \end{aligned}$$

where we arbitrarily pick an $a \in \arg\inf_a \mathbb{E}_{\theta \sim P} L(\theta, a)$ *if there are multiple. Then* $\hat{L}(\theta, P) = L(\theta, f(P))$ *is proper.*

It is possible that by using this trick we remove actions $a \in A$. However, for the purpose of calculating Bayes risks we do not require these actions. From $\hat{L}$, one can define a *regret*

$$D(P, Q) = \mathbb{E}_{\theta \sim P} \hat{L}(\theta, Q) - \mathbb{E}_{\theta \sim P} \hat{L}(\theta, P)$$

which measures how suboptimal the best action is to play against the distribution Q is when played against the distribution P. One does not need knowledge of the function f to construct $\hat{L}$, rather one only needs knowledge of the *Bayes risk*

$$\underline{L}(P) = \inf_a \mathbb{E}_{\theta \sim P} L(\theta, a).$$

From this one can reconstruct $\hat{L}$, and hence L for the purposes of calculating Bayes risks. This is achieved by taking the 1-homogenous extension of $\underline{L}$

$$\begin{aligned} \tilde{\underline{L}} : \mathbb{R}_+^{|\Theta|} &\to \mathbb{R} \\ v &\mapsto \|v\|_1 \underline{L}(\frac{v}{\|v\|_1}) \end{aligned}$$

and differentiating.

Theorem *For a concave $\underline{L} : \mathcal{P}(\Theta) \to \mathbb{R}$*

$$L(\theta, P) = \langle \delta_\theta, \nabla \tilde{\underline{L}}(P) \rangle$$

is a proper loss.

The regret from a proper loss is equal to the *Bregman divergence* defined by $\tilde{\underline{L}}(P)$ and $\underline{L}$,

Theorem *For all concave $\underline{L} : \mathcal{P}(\Theta) \to \mathbb{R}$ and $P, Q \in \mathcal{P}(\Theta)$*

$$D(P, Q) = D_{\underline{L}}(P, Q) = D_{\tilde{\underline{L}}}(P, Q) = \tilde{\underline{L}}(Q) + \langle \nabla \tilde{\underline{L}}(Q), P - Q \rangle - \tilde{\underline{L}}(P).$$

All of these properties show that we only need knowledge of $\underline{L}$ to compute Bayes risks and values.

7.4. Supervised Feature Learning

For a given experiment T and loss function L, supervised feature learning methods aim to minimize the feature gap

$$\Delta \mathcal{V}_L(\pi, \phi, T)$$

The following two lemmas (Banerjee et al., 2005; Tishby et al., 1999) provide means to do this.

Lemma 1 *The feature gap satisfies*

$$\begin{aligned} \Delta \mathcal{V}_L(\pi, \phi, T) &= \mathbb{E}_{x \sim \pi_x} \mathbb{E}_{z \sim \phi(x)} D_{\underline{L}}(T^*(x), (\phi \circ T)^*(z)) \\ &= \mathbb{E}_{x \sim \pi_x} \mathbb{E}_{z \sim \phi(x)} D_{\underline{L}}(P(\Theta|x), P(\Theta|z)) \end{aligned}$$

where $D_{\underline{L}}$ is the regret induced by $\underline{L}$.

Proof

For the proof we use the more familiar probability theory notation with $T^*(x) = P(\Theta|x)$, $(\phi \circ T)^* = P(\Theta|z)$, $\pi_X = P(X)$ and so on. One has

$$D_{\underline{L}}(P, Q) = \underline{L}(Q) + \langle \nabla \underline{L}(Q), P - Q \rangle - \underline{L}(P),\ P, Q \in P(\Theta).$$

and

$$\begin{aligned} \mathcal{V}_L(\pi, T) &= \mathbb{E}_{x \sim P(X)} \underline{L}(P(\Theta|x)) \\ \mathcal{V}_L(\pi, \phi \circ T) &= \mathbb{E}_{x \sim P(Z)} \underline{L}(P(\Theta|z)) \end{aligned}$$

giving

$$\begin{aligned} &\mathbb{E}_{x \sim P(X)} \mathbb{E}_{z \sim P(Z|x)} D_{\underline{L}}(P(\Theta|x), P(\Theta|z)) \\ =&\mathbb{E}_{x \sim P(X)} \mathbb{E}_{z \sim P(Z|x)} \underline{L}(P(\Theta|z)) + \langle \nabla \underline{L}(P(\Theta|z)), P(\Theta|x) - P(\Theta|z) \rangle - \underline{L}(P(\Theta|x)) \\ =&\mathbb{E}_{z \sim P(Z)} \mathbb{E}_{x \sim P(X|z)} \underline{L}(P(\Theta|z)) + \langle \nabla \underline{L}(P(\Theta|z)), P(\Theta|x) - P(\Theta|z) \rangle - \underline{L}(P(\Theta|x)) \\ =&\mathcal{V}_L(\pi, \phi \circ T) - \mathcal{V}_L(\pi, T) + \mathbb{E}_{z \sim P(Z)} \mathbb{E}_{x \sim P(X|z)} \langle \nabla \underline{L}(P(\Theta|z)), P(\Theta|x) - P(\Theta|z) \rangle \end{aligned}$$

Note that $\langle \nabla \underline{L}(P(\Theta|z)), P(\Theta|x) - P(\Theta|z) \rangle$ is affine in $P(\Theta|x)$ and that

$$\mathbb{E}_{x \sim P(X|z)} P(\Theta|x) = P(\Theta|z)$$

meaning

$$\mathbb{E}_{z \sim P(Z)} \mathbb{E}_{x \sim P(X|z)} \langle \nabla \underline{L}(P(\Theta|z)), P(\Theta|x) - P(\Theta|z) \rangle = 0$$

this completes the proof. ■

Lemma 2 *For $\pi_Z \in \mathcal{P}(Z)$, $\hat{U} \in M(Z, \Theta)$*

$$\begin{aligned} &\inf_{\phi} \mathbb{E}_{x \sim \pi_X} \mathbb{E}_{z \sim \phi(x)} D_{\underline{L}}(T^*(x), (\phi \circ T)^*(z)) + \beta I(X, Z) \\ =&\inf_{\phi} \inf_{\hat{\pi}_Z} \inf_{\hat{U}} \mathbb{E}_{x \sim \pi_X} \mathbb{E}_{z \sim \phi(x)} D_{\underline{L}}(T^*(x), \hat{U}(z)) + \mathbb{E}_{x \sim \pi_X} D_{KL}(\phi(x) || \hat{\pi}_Z) \end{aligned}$$

This lemma is proved by the following theorem from Banerjee et al. (2005).

Theorem 4 *For any concave $\underline{L} : X \to \mathbb{R}$ and distribution $P \in \mathcal{P}(X)$*

$$\bar{x} \in \arg\inf_{x} \mathbb{E}_{y \sim x} D_{\underline{L}}(y, x).$$

We can know prove the lemma

Proof Once again we use the more standard notation from probability theory. Firstly

$$I(X; Y) = \mathbb{E}_{x \sim P(X)} D_{KL}(P(Z|x), P(Z)) = \inf_{\hat{P}(Z)} \mathbb{E}_{x \sim P(X)} D_{KL}(P(Z|x), \hat{P}(Z)).$$

since $\mathbb{E}_{x\sim P(X)}P(Z|x) = P(Z)$. Secondly

$$\begin{aligned}\mathbb{E}_{x\sim P(X)}\mathbb{E}_{z\sim P(Z|x)}D_{\underline{L}}(P(\Theta|x), P(\Theta|z)) &= \mathbb{E}_{Z\sim P(Z)}\mathbb{E}_{x\sim P(X|z)}D_{\underline{L}}(P(\Theta|x), P(\Theta|z)) \\ &= \mathbb{E}_{Z\sim P(Z)}\inf_{\hat{P}(\Theta|z)}\mathbb{E}_{x\sim P(X|z)}D_{\underline{L}}(P(\Theta|x), \hat{P}(\Theta|z)) \\ &= \inf_{\hat{P}(\Theta|Z)}\mathbb{E}_{Z\sim P(Z)}\mathbb{E}_{x\sim P(X|z)}D_{\underline{L}}(P(\Theta|x), \hat{P}(\Theta|z))\end{aligned}$$

since $\mathbb{E}_{x\sim P(X|z)}P(\Theta|x) = P(Z|\Theta)$. Combining gives

$$\begin{aligned}&\inf_{P(Z|X)}\mathbb{E}_{x\sim P(X)}\mathbb{E}_{z\sim P(Z|x)}D_{\underline{L}}(P(\Theta|x), P(\Theta|z)) + \beta\mathbb{E}_{x\sim P(X)}D_{KL}(P(Z|x), P(Z)) \\ &= \inf_{P(Z|X)}\inf_{\hat{P}(\Theta|Z)}\inf_{\hat{P}(Z)}\mathbb{E}_{x\sim P(X)}\mathbb{E}_{z\sim P(Z|x)}D_{\underline{L}}(P(\Theta|x), \hat{P}(\Theta|z)) + \beta\mathbb{E}_{x\sim P(X)}D_{KL}(P(Z|x), \hat{P}(Z)).\end{aligned}$$

This completes the proof. ■

These two lemma's yield a family of alternating minimization algorithms that attempt to solve the supervised feature learning problem. The term $\beta I(X;Z)$ can be interpreted as a regularizer that favours ϕ that throw away information about X.

7.5. Variational Divergence

Let P and Q be distributions on a measure space X. The variational divergence has the following equivalent forms

1. $\|P - Q\| = \sup_{\phi:X\to[-1,1]} \mathbb{E}_P\phi - \mathbb{E}_Q\phi$
2. $\|P - Q\| = \inf_C$ such that $\mathbb{E}_P\phi - \mathbb{E}_Q\phi \leq C\|\phi\|_\infty$ for all bounded ϕ.
3. $\|P - Q\| = \int_X |P - Q|$ the l_1 distance between the probability distributions P and Q.
4. $\|P - Q\| = D_f(P, Q) = \int_X f(\frac{dQ}{dP})dP$ the f-divergence from P to Q for $f(x) = |x - 1|$
5. $\|P - Q\| = 2\sup_{A\in\Sigma(X)} P(A) - Q(A), A \subseteq X$

Since $\|\ \|$ is a f-divergence, it also satisfies an information processing theorem (Reid and Williamson, 2011). For any Markov kernel $T : X \rightsquigarrow Y$,

$$\|P - Q\| \geq \|T(P) - T(Q)\|$$

Finally if $\pi \in \mathcal{P}(X)$ and $T, U : X \rightsquigarrow Y$ then

$$\|\pi \otimes T - \pi \otimes U\| = \mathbb{E}_{\theta\sim\pi}\|T(\theta) - U(\theta)\|$$

7.6. Proof that Weighted Deficiency Satisfies the Triangle Inequality

We wish to show that

$$\Delta_\pi(T,U) = \max(\delta_\pi(T,U), \delta_\pi(U,T))$$

provides a metric on experiments on Θ (modulo isomorphism), for priors π that do not assign zero mass to some θ. It should be fairly obvious that Δ_π is both symmetric and non negative. All that is left to prove is the triangle inequality. Fix Markov kernels $T_i : \Theta \rightsquigarrow X_i$ as well as Markov kernels V_1 and V_2 as in the diagram below

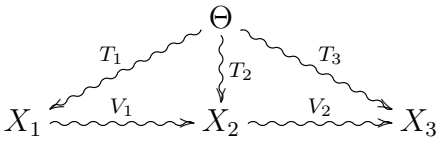

We have $\forall \theta \in \Theta$

$$\begin{aligned}\|T_3(\theta) - V_2 \circ V_1 \circ T_1(\theta)\| &\leq \|T_3(\theta) - V_2 \circ T_2(\theta)\| + \|V_2 \circ T_2(\theta) - V_2 \circ V_1 \circ T_1(\theta)\| \\ &\leq \|T_3(\theta) - V_2 \circ T_2(\theta)\| + \|T_2(\theta) - V_1 \circ T_1(\theta)\|\end{aligned}$$

where we have used the fact the variational divergence is a metric and that it satisfies an information processing theorem. Averaging with respect to the prior and taking infimums over V_1 and V_2 yields

$$\delta_\pi(T_1||T_3) \leq \delta_\pi(T_1||T_2) + \delta_\pi(T_2||T_3).$$

Going in the opposite direction and taking maximums yields

$$\Delta_\pi(T_1,T_3) \leq \Delta_\pi(T_1,T_2) + \Delta_\pi(T_2,T_3)$$

the desired result. Even if π does give zero mass to some θ, meaning Δ_π may not be a metric, the triangle inequality still applies.

7.7. Proof of the Information Processing Theorem

Theorem (Blackwell-Sherman-Stein) *$T|U$ if and only if $\mathcal{V}_L(\pi,T) \leq \mathcal{V}_L(\pi,U)$ for all loss functions L and priors π.*

Proof

By definition we have

$$\begin{aligned}\mathcal{V}_L(\pi,T) &= \inf_{D \in M(\Theta,A)_T} R_L(\pi,D) \\ &= \inf_{D \in M(\Theta,A)_T} f(D)\end{aligned}$$

Where $f : M(\Theta,A) \to \mathbb{R}$. By property (1) of factoring through we have $M(\Theta,A)_T \supseteq M(\Theta,A)_U$ giving

$$\begin{aligned}\mathcal{V}_L(\pi,T) &= \inf_{D \in M(\Theta,A)_T} f(D) \\ &\leq \inf_{D \in M(\Theta,A)_U} f(D) \\ &= \mathcal{V}_L(\pi,U)\end{aligned}$$

■

7.8. Proof of the Blackwell-Sherman-Stein theorem

Theorem (Blackwell-Sherman-Stein) *$T|U$ if and only if $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U)$ for all loss functions L and priors π.*

Here we prove the converse to the information processing theorem.

Proof

The condition that $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U)$ for all loss functions and priors is equivalent to

$$\forall L,\ \forall D_U \in M(\Theta, A)_U,\ \exists D_T \in M(\Theta, A)_T \text{ such that } R_L(\delta_\theta, D_Y) \leq R_L(\delta_\theta, d_Y),\ \forall \theta$$

in words, for any loss and any decision rule based on U there is one based on T that is better.

To see this note for each $\theta_0 \in \Theta$ we can take loss functions that only care about that particular θ_0, ie $L(\theta, a) = 0$ if $\theta \neq \theta_0$. Now fix A. We have that $R_L(\delta_\theta, -) : M(\Theta, A) \to \mathbb{R}$ is linear. Furthermore $M(\Theta, A)_{P_T}$ and $M(\Theta, A)_U$ are closed convex subsets of $M(\Theta, A)$. As we vary L and θ, $R_L(\delta_\theta, -)$ gives the entire dual space of $M(\Theta, A)$, ie by taking $L(\theta_0, a_0) = 1$ and zero for all other θ and a. By the correspondence between risk functions and the dual of $M(\Theta, A)$ means the above is equivalent to

$$\inf_{D \in M(\Theta, A)_T} \langle \alpha, D \rangle \leq \inf_{D \in M(\Theta, A)_U} \langle \alpha, D \rangle,\ \forall \alpha \in M(\Theta, A)^*$$

meaning that the support function of $M(\Theta, A)_T$ is less than or equal to the support function of $M(\Theta, A)_T$. Hence

$$M(\Theta, A)_U \supseteq M(\Theta, A)_T$$

Note A was arbitrary. Now let $A = Y$, and note that $U \in M(\Theta, Y)_U$. Hence by the above inclusion, there exists a $U/T : X \to Y$ with $U = (U/T) \circ T$

■

7.9. Proof of the Randomization theorem

Theorem (Randomization) *Fix $\epsilon > 0$, T, U and π. $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U) + \epsilon \|L\|_\infty$ if and only if $\delta_\pi(T, U) \leq \epsilon$ for all loss functions L.*

Firstly the if direction.

Proof (Forward implication)

If $\delta(T||U) \leq \epsilon$ then there exists a Markov kernel V such that

$$\mathbb{E}_{\theta \sim \pi} \|U(\theta) - V \circ T(\theta)\| \leq \epsilon$$

Fix a decision rule $D_U \in M(\Theta, A)_U$. Using V gives a decision rule $D_T \in M(\Theta, A)_T$ by composition, $D_T = (D_U/U) \circ V \circ T$. See the diagram below for a better explanation.

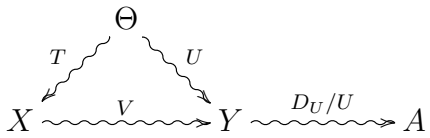

By the properties of the variational divergence one has

$$\begin{aligned} R_L(\pi, D_T) - R_L(\pi, D_U) &= \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{a\sim D_T(\theta)}L(\theta, a) - \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{a\sim D_U(\theta)}L(\theta, a) \\ &\leq \mathbb{E}_{\theta\sim\pi}\|D_U(\theta) - D_T(\theta)\|\|L\|_\infty \\ &= \mathbb{E}_{\theta\sim\pi}\|(D_U/U) \circ U(\theta) - ((D_U/U) \circ V \circ T)(\theta)\|\|L\|_\infty \\ &\leq \mathbb{E}_{\theta\sim\pi}\|U(\theta) - (V \circ T)(\theta)\|\|L\|_\infty \\ &\leq \epsilon\|L\|_\infty \end{aligned}$$

where the first line is the definition of risk, the second follows from one definition of the variational divergence and the third follows from the fact the variational divergence is itself an f-divergence and as such satisfies an information processing theorem. Note this holds for all $D_U \in M(\Theta, A)_U$.

Taking infimums over D_T and D_U, one has $\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U) + \epsilon\|L\|_\infty$ for all loss functions. ■

We now prove the converse.

Proof (Converse)

Fix A, a decision rule $D_U \in M(\Theta, A)_U$ and a prior π and define a function

$$f(L, D_T) = R_L(\pi, D_T) - R_L(\pi, D_U) - \epsilon\|L\|_\infty$$

that takes a decision rule $D_T \in M(\Theta, A)_T$ and a loss L and returns the difference in Bayes risks of D_T and D_U subtracted by a term that penalizes losses with high magnitude. Note that f is a linear in D_T and concave in L. By the conditions in the theorem one has

$$\sup_L \inf_{D_T} f(L, D_T) \leq 0.$$

By the minimax theorem (Komiya, 1988) , there exists a saddle point (L^*, D_T^*) with

$$f(L^*, D_T^*) = \sup_L \inf_{D_T} f(L, D_T) = \inf_{D_T} \sup_L f(L, D_T).$$

Hence $f(L, D_T^*) \leq 0,\ \forall L$, meaning

$$R_L(\pi, D_T^*) - R_L(\pi, D_U) \leq \epsilon\|L\|_\infty\ \forall L.$$

This implies that

$$\mathbb{E}_{\theta\sim\pi}\mathbb{E}_{a\sim D_T^*(\theta)}L(\theta, a) - \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{a\sim D_U(\theta)}L(\theta, a) \leq \epsilon\|L\|_\infty\ \forall L$$

meaning

$$\|D_T^* \otimes \pi - D_U \otimes \pi\| = \mathbb{E}_{\theta\sim\pi}\|D_T(\theta) - D_U(\theta)\| \leq \epsilon$$

As D_U was arbitrary, we have for all $D_U \in M(\Theta, A)_U$ there exists a $D_T \in M(\Theta, A)_T$ with $\mathbb{E}_{\theta\sim\pi}\|D_T(\theta) - D_U(\theta)\| \leq \epsilon$.

Once again, take $Y = A$ and note that $U \in M(\Theta, Y)_U$.

■

7.10. Proof of Theorem 1

Theorem *For all experiments U, T and all priors π*

$$\Delta_\pi(U, T) = \sup_L \frac{|\mathcal{V}_L(\pi, U) - \mathcal{V}_L(\pi, T)|}{\|L\|_\infty}$$

For the proof we require the following very simple lemma.

Lemma 3 *For $x, y \in \mathbb{R}$ if $\forall \epsilon \in \mathbb{R}$ we have $x \leq \epsilon \Leftrightarrow y \leq \epsilon$ then $x = y$.*

Proof Suppose that $x \neq y$ and without loss of generality assume that $x \leq y$. Set $\epsilon = \frac{x+y}{2}$. Then $x \leq \epsilon$ and $y \geq \epsilon$, a contradiction. ■

Now the theorem.

Proof If $\Delta_\pi(U, T) \leq \epsilon$ then $\delta_\pi(U, T) \leq \epsilon$ and $\delta_\pi(T, U) \leq \epsilon$. By the randomization theorem

$$|\mathcal{V}_L(\pi, U) - \mathcal{V}_L(\pi, T)| \leq \epsilon\|L\|_\infty, \ \forall L$$

hence

$$\sup_L \frac{|\mathcal{V}_L(\pi, U) - \mathcal{V}_L(\pi, T)|}{\|L\|_\infty} \leq \epsilon.$$

Conversely if $\sup_L \frac{|\mathcal{V}_L(\pi,U)-\mathcal{V}_L(\pi,T)|}{\|L\|_\infty} \leq \epsilon$ then for all L

$$\mathcal{V}_L(\pi, T) \leq \mathcal{V}_L(\pi, U) + \epsilon\|L\|_\infty$$

and

$$\mathcal{V}_L(\pi, U) \leq \mathcal{V}_L(\pi, T) + \epsilon\|L\|_\infty$$

which by the randomization theorem gives $\Delta_\pi(U, T) \leq \epsilon$. Combining these two facts and the lemma completes the proof. ■

7.11. Proof of Theorem 2

Theorem *For all experiments $T : \Theta \rightsquigarrow X$, for all measure spaces Z and for all feature maps $\phi : X \rightsquigarrow Z$*

$$\Delta_\pi(T, \phi \circ T) \leq \Delta_{T(\pi)}(id_X, \phi).$$

To prove this theorem we note that from the randomization theorem we have (Torgersen, 1991)

$$\Delta_\pi(T, U) = \sup_L \frac{|\mathcal{V}_L(\pi, T) - \mathcal{V}_L(\pi, U)|}{\|L\|_\infty}.$$

By standard manipulations from decision theory one has

$$\begin{aligned}
\mathcal{V}_L(\pi, T) &= \inf_{d:X \rightsquigarrow A} \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{x\sim T(\theta)}\mathbb{E}_{a\sim d(x)}L(\theta, a) \\
&= \inf_{d:X \rightsquigarrow A} \mathbb{E}_{x\sim\pi_X}\mathbb{E}_{\theta\sim T^*(x)}\mathbb{E}_{a\sim d(x)}L(\theta, a) \\
&= \mathbb{E}_{x\sim\pi_X} \inf_{a\in A} \mathbb{E}_{\theta\sim T^*(x)}L(\theta, a) \\
&= \mathbb{E}_{x\sim\pi_X}\underline{L}(T^*(x))
\end{aligned}$$

giving

$$\Delta_\pi(T, U) = \sup_L \frac{|\mathbb{E}_{x\sim T(\pi)}\underline{L}(T^*(x)) - \mathbb{E}_{y\sim U(\pi)}\underline{L}(U^*(x))|}{\|L\|_\infty}.$$

Due to the correspondence between loss functions and their Bayes risks (Reid and Williamson, 2011; Dawid, 2007; Grünwald and Dawid, 2004; Garcia-Garcia and Williamson, 2012) (In particular one can recover the loss from its Bayes Risk), and the fact that Bayes risks $\underline{L} : \mathcal{P}(\Theta) \to \mathbb{R}$ are "attached directly to Θ", in feature learning we may wish to move them to objects that are attached to X. We are inspired by the diagram below.

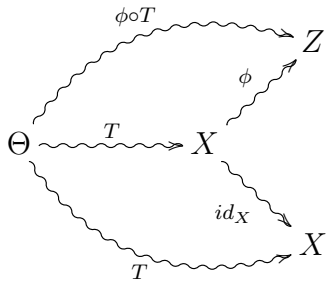

This is achieved by the following lemma.

Lemma 4 *Let $\mathcal{L}(\Theta)$ denote the set of all Bayes risks on Θ (proper concave functions). Fix a Markov kernel $T : \Theta \rightsquigarrow X$ and a prior π on Θ. Then the Markov kernel $T^* : \Theta \rightsquigarrow \Theta$ provided by Bayes rule gives a function*

$$\begin{aligned}
&\underline{\mathcal{L}}[T] : \mathcal{L}(\Theta) \to \mathcal{L}(X) \\
&\underline{L} \mapsto \underline{L} \circ T^*
\end{aligned}$$

furthermore for a chain of Markov kernels

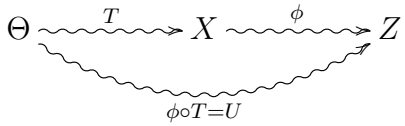

we have that $\mathbb{E}_{z\sim U(\pi)}\underline{L}(U^*(y)) = \mathbb{E}_{z\sim U(\pi)}\underline{\mathcal{L}}[T](\underline{L})(\phi^*(y))$

Proof

It should be fairly obvious that $\underline{\mathcal{L}}[T](\underline{L})$ is concave as it is the composition of a linear function and a concave function. Furthermore

$$\begin{aligned}\mathbb{E}_{z\sim U(\pi)}\underline{\mathcal{L}}[T](\underline{L})(\phi^*(z)) &= \mathbb{E}_{z\sim U(\pi)}\underline{L}(T^*\circ\phi^*(z)) \\ &= \mathbb{E}_{z\sim U(\pi)}\underline{L}(U^*(z))\end{aligned}$$

■

we can now prove theorem 2

Proof

$$\Delta_\pi(T,\phi\circ T) = \sup_{\underline{L}\in\mathcal{L}(\Theta)} \frac{|\mathbb{E}_{x\sim T(\pi)}\underline{L}(T^*(x)) - \mathbb{E}_{y\sim\phi\circ T(\pi)}\underline{L}(T^*\circ\phi^*(x))|}{\|L\|_\infty} \tag{1}$$

$$= \sup_{\underline{L}\in\mathcal{L}(\Theta)} \frac{|\mathbb{E}_{x\sim T(\pi)}\underline{\mathcal{L}}[T](\underline{L})(\delta_x) - \mathbb{E}_{y\sim\phi\circ T(\pi)}\underline{\mathcal{L}}[T](\underline{L})(\phi^*(y))|}{\|\underline{\mathcal{L}}[T](L)\|_\pi} \tag{2}$$

$$\leq \sup_{\underline{L}\in\mathcal{L}(X)} \frac{|\mathbb{E}_{x\sim T(\pi)}\underline{L}(\delta_x) - \mathbb{E}_{y\sim\phi\circ T(\pi)}\underline{L}(\phi^*(x))|}{\|L\|_\infty} \tag{3}$$

$$= \Delta_{T(\pi)}(id_X,\phi)$$

■

7.12. Proof of Theorem 3

Theorem *Fix a prior* π_X *and an* $\epsilon > 0$. $\phi : X \rightsquigarrow Z$ *constitutes generic features of quality* ϵ *for* X *if and only if for the reconstruction problem* (X, Z, ϕ, X, L_{01})*, one can find a decision rule* $d : Z \rightsquigarrow X$ *with*

$$R_{L_{01}}(\pi_X, d\circ\phi) \leq \frac{\epsilon}{2}.$$

First the if.

Proof (Forward Implication) Let $T : \Theta \rightsquigarrow X$ be an experiment, L a loss and π a prior on Θ with $T(\pi) = \pi_X$. For any $D_T \in M(\Theta, A)_T$ consider

$$D_{\phi\circ T} = (D_T/T)\circ d\circ\phi\circ T \in M(\Theta,A)_{\phi\circ T}.$$

We wish to calculate

$$R_L(\pi, D_{\phi\circ T}) = \mathbb{E}_{\theta\sim\pi}\mathbb{E}_{x\sim T(\theta)}\mathbb{E}_{x'\sim(d\circ\phi)(x)}\mathbb{E}_{a\sim(D_T/T)(x')}L(\Theta, x).$$

To calculate this risk we consider two cases. Either $x = x'$ with probability $1 - \frac{\epsilon}{2}$, giving risk $R_L(\pi, D_T)$, or $x' \neq x$ with probability $\frac{\epsilon}{2}$ with risk bounded above by $2\|L\|_\infty$. This upper bound occurs when the action chosen by d after seeing x is the best for θ and the action chosen by d after seeing x' is the worst for θ. Combining these two gives

$$R_L(\pi, D_{\phi\circ T}) \leq (1 - \frac{\epsilon}{2})R_L(\pi, D_T) + \frac{\epsilon}{2}2\|L\|_\infty \leq R_L(\pi, D_T) + \epsilon\|L\|_\infty.$$

Note that T and L are arbitrary. Taking infimums over D_T and $D_{\phi\circ T}$ gives

$$\mathcal{V}_L(\pi, \phi \circ T) \leq \mathcal{V}_L(\pi, T) + \epsilon\|L\|_\infty,\ \forall L,\ \forall T,\ \forall \pi\ st\ T(\pi) = \pi_X.$$

■

Now the converse.

Proof (Converse)

If

$$\mathcal{V}_L(\pi, \phi \circ T) \leq \mathcal{V}_L(\pi, T) + \epsilon\|L\|_\infty,\ \forall L,\ \forall T,\ \forall \pi\ st\ T(\pi) = \pi_X$$

then

$$\Delta_\pi(\phi \circ T, T) \leq \epsilon,\ \forall T,\ \forall \pi\ st\ T(\pi) = \pi_X$$

in particular by taking $T = \mathrm{id}_X$ and $\pi = \pi_X$ we have

$$\Delta_{\pi_X}(\phi, \mathrm{id}_x) \leq \epsilon.$$

By the definition of weighted deficiency and as $\mathrm{id}_X\ |\phi$, $\Delta_{\pi_X}(\phi, \mathrm{id}_x) = \delta_\pi(\phi, \mathrm{id}_X)$. Therefore there exists a markov kernel $d : Z \rightsquigarrow X$ with

$$\mathbb{E}_{x\sim\pi_X}\|\delta_x - (d \circ \phi)(x)\| = 2\mathbb{E}_{x\sim\pi_X}\mathbb{E}_{x'\sim(d\circ\phi)(x)}\mathbb{1}(x' = x) \leq \epsilon$$

Hence $R_{L_{01}}(\pi_X, d \circ \phi) \leq \frac{\epsilon}{2}$.

■

References

Arindam Banerjee, Srujana Merugu, Inderjit S. Dhillon, and Joydeep Ghosh. Clustering with Bregman divergences. *The Journal of Machine Learning Research*, 6:1705–1749, 2005.

Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

Anthony J. Bell and Terrence J. Sejnowski. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 7(6):1129–1159, 1995.

Yoshua Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1): 1–127, 2009.

David Blackwell. Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102, 1951.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2): 265–272, 1953.

Thomas M. Cover and Jay A. Thomas. *Elements of Information Theory*. Wiley, 2012.

A. Phillip Dawid. The geometry of proper scoring rules. *Annals of the Institute of Statistical Mathematics*, (April 2006):77–93, 2007.

Morris H. DeGroot. Uncertainty, information, and sequential experiments. *The Annals of Mathematical Statistics*, 33(2):404–419, 1962.

Thomas Shelburne Ferguson. *Mathematical statistics: A decision theoretic approach*. Academic Press New York, 1967.

Dario Garcia-Garcia and Robert C. Williamson. Divergences and Risks for Multiclass Experiments. In *Conference on Learning Theory (JMLR: W&CP)*, volume 23, 2012.

Peter D. Grünwald and A. Philip Dawid. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *The Annals of Statistics*, 32(4):1367–1433, 2004.

Martin Hellman and Josef Raviv. Probability of error, equivocation, and the Chernoff bound. *IEEE Transactions on Information Theory*, 16(4):368–372, 1970.

Geoffrey E. Hinton and Ruslan R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

Peter J. Huber. *Robust Statistics*. Wiley Series in Probability and Statistics. Wiley, 2011.

Hidetoshi Komiya. Elementary proof for Sion's minimax theorem. *Kodai mathematical journal*, 11(1):5–7, 1988. ISSN 0386-5991.

Lucien Lecam. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, 35(4): 1419–1455, 1964.

Lucien Lecam. On the Information Contained in Additional Observations. *The Annals of Statistics*, 2(4): 630–649, 1974.

Lucien Lecam. *Asymptotic Methods in Statistical Decision Theory*. Springer London, 2011.

Yann LeCun. Learning Representations: A Challenge for Learning Theory, Plenary talk COLT 2013, 2013.

Honglak Lee, Alexis Battle, Rajat Raina, and Andrew Ng. Efficient sparse coding algorithms. In *Advances in neural information processing systems*, pages 801–808, 2006.

Friedrich Liese. phi-divergences, sufficiency, Bayes sufficiency, and deficiency. *Kybernetika*, 48(4):690–713, 2012.

Ralph Linsker. *An application of the principle of maximum information preservation to linear systems*. NIPS, 1989.

Bruno A. Olshausen and David J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research*, 37(23):3311–3325, 1997.

Matthew Parry, A Philip Dawid, and Steffen Lauritzen. Proper local scoring rules. *The Annals of Statistics*, 40(1):561–592, 2012.

Mark D. Reid and Robert C. Williamson. Information, divergence and risk for binary experiments. *The Journal of Machine Learning Research*, 12:731–817, 2011.

Vin D. Silva and Joshua B. Tenenbaum. Global versus local methods in nonlinear dimensionality reduction. In *Advances in neural information processing systems*, pages 705–712, 2002.

Joshua B. Tenenbaum, Vin De Silva, and John C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

Naftali Tishby, Fernando C. Pereira, and William Bialek. The information bottleneck method. In *Proc. of Allerton Conf. on Communication, Control and Computing*, volume physics/00, pages 368–377, 1999.

Erik Torgersen. *Comparison of Statistical Experiments*. Cambridge University Press, 1991.

Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th International Conference on Machine Learning*, pages 1096–1103. ACM, 2008.

Moshe Zakai and Jacob Ziv. A generalization of the rate-distortion theory and application. *Information Theory, New Trends and Open Problems*, pages 87–123, 1975.

Jacob Ziv and Moshe Zakai. On functionals satisfying a data-processing theorem. *IEEE Transactions on Information Theory*, I(3):770–772, 1973.